

RECEIVED
CENTRAL FAX CENTER

FEB 08 2006

**Yee &
Associates, P.C.**

4100 Alpha Road
Suite 1100
Dallas, Texas 75244

Main No. (972) 385-8777
Facsimile (972) 385-7766

Facsimile Cover Sheet

To: Commissioner for Patents for Examiner Sathyanaraya R. Pannala Group Art Unit 2164	Facsimile No.: 571/273-8300
From: Lourdes Perez Legal Assistant to Cathrine K. Kinslow	No. of Pages Including Cover Sheet: 25
Message: Transmitted herewith: <ul style="list-style-type: none">• Transmittal Document; and• Reply Brief.	
Re: Docket No: AUS920010286US1	Serial No.: 09/895,231
Date: Wednesday, February 08, 2006	
Please contact us at (972) 385-8777 if you do not receive all pages indicated above or experience any difficulty in receiving this facsimile.	<i>This Facsimile is intended only for the use of the addressee and, if the addressee is a client or their agent, contains privileged and confidential information. If you are not the intended recipient of this facsimile, you have received this facsimile inadvertently and in error. Any review, dissemination, distribution, or copying is strictly prohibited. If you received this facsimile in error, please notify us by telephone and return the facsimile to us immediately.</i>

**PLEASE CONFIRM RECEIPT OF THIS TRANSMISSION BY
FAXING A CONFIRMATION TO 972-385-7766.**

RECEIVED
CENTRAL FAX CENTER

FEB 08 2006

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Koelle et al.

Serial No.: 09/895,231

Filed: June 29, 2001

For: Decentralized, Self-Regulating System for Automatically Discovering Optimal Configurations in a Failure-Rich Environment

35525

PATENT TRADEMARK OFFICE
CUSTOMER NUMBER

Group Art Unit: 2164

Examiner: Pannala, Sathyanaraya R.

Attorney Docket No.: AUS920010286US1

§ Certificate of Transmission Under 37 C.F.R. § 1.8(a)
 § I hereby certify this correspondence is being transmitted via
 § facsimile to the Commissioner for Patents, P.O. Box 1450,
 § Alexandria, VA 22313-1450, facsimile number (571) 273-8300,
 § on February 8, 2006.

By:

Louides Perez

TRANSMITTAL DOCUMENT

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir

TRANSMITTED HEREWITH:

- **Reply Brief (37 C.F.R. 41.41).**

No fees are believed to be required. If, however, any fees are required, I authorize the Commissioner to charge these fees which may be required to IBM Corporation Deposit Account No. 09-0447.

Respectfully submitted,

Cathrine K. Kinslow
Registration No. 51,886

Duke W. Yee
Registration No. 34,285

YEE & ASSOCIATES, P.C.
P.O. Box 802333
Dallas, Texas 75380
(972) 385-8777
ATTORNEYS FOR APPLICANTS

**RECEIVED
CENTRAL FAX CENTER**

FEB 08 2006

Docket No. AUS920010286US1

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Koelle et al.

25

Serial No. 09/895,231

223

Group Art Unit: 2164

2025

Filed: June 29, 2001

22

Examiner: Pannala, Sathyanaraya R.

25

For: Decentralized, Self-Regulating System for Automatically Discovering Optimal Configurations in a Failure-Rich Environment

22

20


32

20

2

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Certificate of Transmission Under 37 C.F.R. § 1.8(a)

I hereby certify this correspondence is being transmitted via facsimile to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, facsimile number (571) 273-8300, on February 8, 2006. 

By:

Lourdes Perez

REPLY BRIEF (37 C.F.R. 41.41)

This brief is submitted in response to the Examiner's Answer mailed on December 8, 2005.

No fees are believed to be required to file a Reply Brief. Any required petition for extension of time for filing this brief and fees therefore, are dealt with in the accompanying TRANSMITTAL OF REPLY BRIEF.

Reply Brief Page 1 of 23
Koelle et al. - 09/895,231

RESPONSE TO EXAMINER'S ANSWER**IA. 35 U.S.C. § 102, Alleged Anticipation of Claims 1-5, 10, 12-16, and 21-25**

With regard to claim 1, Appellants respectfully disagree with the Examiner's assertions that the *Kenner* reference teaches "organizing the set of data into a plurality of related sets of data", "assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria", and "transferring management of the related set of data managed by the failed service to another service within the distributed set of services in response to failure of a service within the distributed set of services".

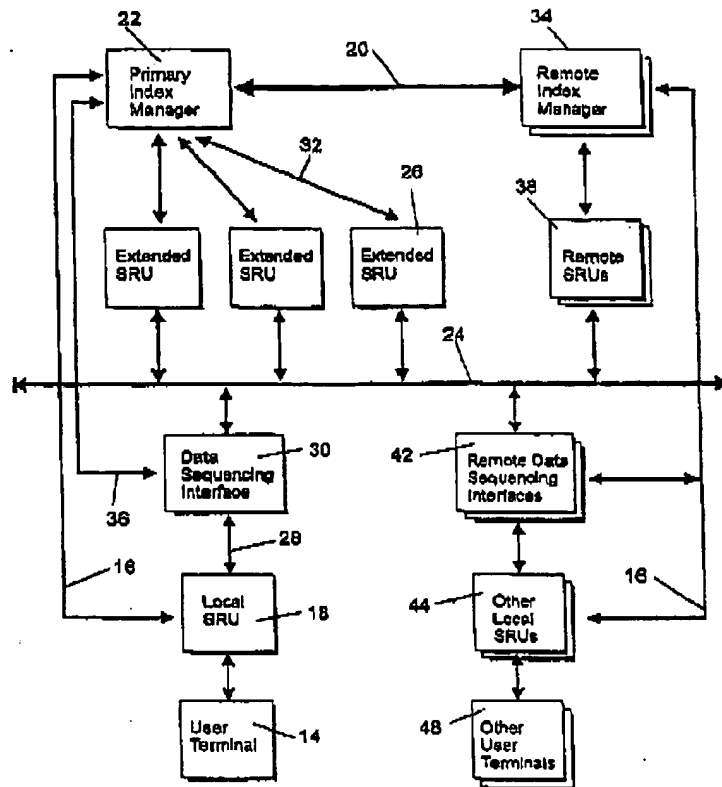
With respect to "organizing the set of data into a plurality of related sets of data," the Examiner states:

...Kenner do teach as the SRU command logic sees to the duplication of popular videos on alternate SRUs 26. It also places copies of video segments on SRUs geographically closer to the user most interest in those videos. Duplication of data is done during the non-peak periods of the system (Fig. 1, col. 8, lines 27-32). Kenner also teaches as the system provides a plurality of extended SRUs, each of which communicates with PIM and date Sequencing Interface (DSI). This provides a flexible, high capacity, high throughput system which can be readily extended as needed, and can provide for efficient distribution and backup of video clips and other data on the system (Fig. 1, col. 14, lines 56-63). Further, Appellant argued as "Kenner fails to teach assignment of data sets based on optimization criteria." Again the Examiner respectfully disagrees with the Appellant because the current invention on page 14, lines 1-3 stated as "an optimal service may be used for load balancing data between services." Kenner do teach load balancing as the SRU under-run counter parameter identifies the location of "over-accessed" SRUs, audio-visual data will be moved or copied from these heavily loaded SRUs to more lightly loaded SRUs based on their under-run levels, in an effort to distribute or flatten SRU demand (col. 14, lines 33-39).

Examiner's Answer dated December 8, 2005, pp. 26-27.

Figure 1 and column 8, lines 27-32 of *Kenner* are reproduced below:

FIGURE 1



Kenner, Figure 1.

The "SRU Command logic" sees to the duplication of popular videos on alternate SRUs 26. It also places copies of video segments on SRUs geographically closer to the users most interested in those videos. The goal is not duplicate data onto SRUs 26 where the number of frequently downloaded videos ("FDVs") is already high (above a predetermined value).

Kenner, col. 8, lines 27-32.

Figure 1 shows a video clip storage and retrieval system, comprising a plurality of extended storage and retrieval units (SRUs) 26. As discussed in the passage above, *Kenner* discloses that a popular video clip may be duplicated and stored on alternate SRUs 26. A duplicate copy of a video clip is placed on the alternate SRU closest to users most interested in the clip. Thus, *Kenner* discloses identifying users' interest in a video clip, duplicating the clip, and placing the duplicated clip on the SRU that is geographically closest to those users.

However, *Kenner* does not teach organizing video clips into related sets. As shown in the passage above, *Kenner* merely duplicates individual video clips to whichever SRU needs them. No organization of the clips into related sets is mentioned in *Kenner*, because *Kenner* is merely concerned with selecting individual clips that are accessed most frequently, and then providing each individual clip to the closest SRU where the clip was used most. Consequently, *Kenner* does not teach organizing the set of data into a plurality of related sets of data as recited in claim 1 of the present invention.

In addition, *Kenner* does not teach "assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria", nor does *Kenner* mention "transferring management of the related set of data managed by the failed service to another service". The Examiner alleges these features are found in the following cited section of *Kenner*:

Whenever an SRU fails to deliver the requested video clip, the DSI 30 increments the SRU under-run counter for that SRU and eventually communicates this information to the PIM 22. If the SRU under-run count exceeds a predetermined threshold value (communicated to the DSI 30 upon creation), the PIM 22 directs further requests away from this affected SRU by having the DSI 30 query alternate SRUs for the video clip information. In the event that the video clip is only stored at this location, then a delay will be encountered as the DSI 30 waits for the video information to be downloaded. The PIM 22 will also direct that the number of FDVs be decremented for this affected extended SRU 26.

In addition, since the SRU under-run count parameter identifies the location of "over-accessed" SRUs, audio-visual data will be moved or copied from these heavily loaded SRUs to more lightly loaded SRUs (based on their under-run levels), in an effort to distribute or flatten SRU demand. This load management process will occur during off-peak hours.

Kenner, col. 14, lines 20-39.

The passage above discloses that a count parameter is used to identify over-accessed SRUs in the storage retrieval system, and that copies of the video clips may be distributed from more heavily loaded SRUs to lightly loaded SRUs. The passage also discloses that a data sequencing interface (DSI) is used to increment the SRU counter and communicates this information to the primary index managers (PIM). If the count for the SRU exceeds a threshold, the PIM directs requests for video clips on that SRU to an alternate SRU.

However, the passage above does not teach assigning or transferring management of a

related set of video clips to another service. *Kenner* merely teaches that if an SRU fails to deliver the requested individual video clip, the request is directed to another SRU which already manages the video clip. *Kenner* does not teach that management of a related set of data is assigned, but rather relies on individual video clips being duplicated. In addition, since *Kenner* teaches that a video clip may be handled by another SRU only if that other SRU already manages the same video clip, no change in the assignment of or transfer of management occurs. Consequently, *Kenner* does not teach assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria, or transferring management of the related set of data managed by the failed service to another service within the distributed set of services in response to failure of a service within the distributed set of services, as recited in claim 1 of the present invention.

In view of the above, *Kenner* does not anticipate claim 1. Independent claims 10, 12, and 21 recite subject matter addressed above with respect to claim 1 and are allowable for similar reasons. Since claims 2-5, 13-16, and 22-25 depend from claims 1, 12, and 21, the same distinctions between *Kenner* and the invention recited in claims 1, 12, and 21 apply for these claims. Additionally, claims 2-5, 13-16, and 22-25 recite other additional combinations of features not taught by the reference.

Therefore, Appellants respectfully request that the rejection of claims 1-5, 10, 12-16, and 21-25 under 35 U.S.C. § 102(e) not be sustained.

IA(1). 35 U.S.C. § 102, Alleged Anticipation of Claims 2, 13, and 22

With respect to claim 2, the Examiner's Answer states:

Claims 2, 13 and 22 are dependent on independent claim 1, 12 and 21 respectively. All three claims are basically claiming the same. Claim 2 stated as "The method as recited in claim 1, wherein the optimization criteria is based on location of the service within the distributed set of services." *Kenner* teaches as the Audio-visual data index database is searched to determine most frequently accessed SRUs in comparison to predetermined value. Those extended SRUs are selected for the duplication or transferal. The selected SRUs are evaluated to whether they can accept duplicate copy of the video clip. If so, the FDV is duplicated on the identified extended SRU 26 (Fig. 1, col. 8, lines 38-47). The PIM places copies of video segments on SRUs geographically closer to the users most interested in those videos (Fig. 1, col. 8, lines 28-30).

Appellant argues on page 18, as "However, the Final Office Action proffers no explanation as to how duplicating video clips based on the most

frequently requested is somehow equivalent to assigning management of data sets to by data services based on location of service." Examiner respectfully disagrees with the Appellant because, management of data sets by data services and duplicate copies of video clips placing on SRUs geographically closer to the users most interested in those videos. In this case, the information needed is placed closer so that the user can access easily and quicker than getting from a central location through communication network.

Examiner's Answer dated December 8, 2005, pp. 28-29.

Kenner does not teach the feature of "wherein the optimization criteria is based on location of the service within the distributed set of services". The Examiner alleges this feature is found in column 8, lines 38-47 and 28-30 of *Kenner*, which are reproduced below:

This determination is made by searching through the "Audio-Visual Data Index" database (described below) to identify the video clips that have been accessed most frequently. From this video subset, videos are selected for transferal or duplication based on where the video was used most. If the FDV was transferred principally from DSIs 30 created by the PIM 22, extended SRUs 26 located within the same computer are evaluated to determine whether that extended SRU 26 can accept a duplicate copy of the video clip. If so, the FDV is duplicated on the identified extended SRU 26.

Kenner, col. 8, lines 38-47.

The "SRU Command logic" sees to the duplication of popular videos on alternate SRUs 26. It also places copies of video segments on SRUs geographically closer to the users most interested in those videos.

Kenner, col. 8, lines 28-30.

The passages above disclose identifying video clips by how often they are accessed ("accessed most frequently"), and that duplication of the identified most frequently accessed video clips are duplicated on SRUs that are closest to users that are accessing the video clips. Thus, *Kenner's* determination of whether a video clip is duplicated to another SRU is based on how often the individual video clip is accessed. Thus, *Kenner* teaches that the criteria for duplicating an individual video clip on another SRU is based on how often the video clip is accessed, rather than based on the location of the services. Thus *Kenner*, in fact, does not teach assigning management of data sets by data services based on location of the data services. Claims 13 and 22 recite subject matter addressed above with respect to claim 2 and are allowable for similar reasons.

IA(2). 35 U.S.C. § 102, Alleged Anticipation of Claims 3, 14, and 23

With respect to claim 3, the Examiner's Answer states:

Claims 3, 14 and 23 are dependent on independent claims 1, 12 and 21 respectively. All three claims are basically claiming the same. Claim 3 stated as "The method as recited in claim 1, detecting the failed service by a set of remaining services within the distributed set of services and examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service" as whenever an SRU fails to deliver the requested video clip, the DSI 30 increments the SRU under-run counter for that SRU and eventually communicates this information to the PIM 22. If the SRU under-run count exceeds a predetermined threshold value the PIM 32 directs further requests away from this affected SRU by the DSI 30 query alternate SRUs for the video clip information (Fig. 1, col. 14, lines 20-28).

Appellant argues on page 19, as "Appellants submit that *Kenner*, in fact, does not teach or fairly suggest a distributed set of data services wherein remaining data service examine data sets managed by a failed data service." Examiner respectfully disagrees with the Appellant argument because of misinterpretation of the reference. *Kenner* teaches as whenever an SRU fails to deliver the requested video clip, the DSI 30 increments the SRU under-run counter for that SRU and eventually communicates this information to the PIM 22. If the SRU under-run count exceeds a predetermined threshold value the PIM 32 directs further requests away from this affected SRU by the DSI 30 query alternate SRUs for the video clip information (Fig. 1, col. 14, lines 20-28).

Examiner's Answer dated December 8, 2005, pp. 29-30.

Kenner does not teach the feature of "examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service". The Examiner points to column 14, lines 20-28 of *Kenner* (reproduced above) as teaching these features. However, as shown above, the passage discloses that a data sequencing interface (DSI) is used to increment the SRU counter and communicates this information to the primary index managers (PIM), which directs requests for video clips on that SRU to an alternate SRU if the count for that SRU exceeds a threshold. However, directing a request for an individual video clip to another SRU based on the SRU counter is not the same as the related set of data managed by a failed service examined by the set of remaining services in the system. There is no mention in *Kenner* that remaining data services examine the related data sets managed by a failed data service. Instead, *Kenner* merely teaches that if an SRU fails to respond to a request for an individual video clip, an alternate SRU having a copy of the video clip may receive the request. Consequently, *Kenner* does not anticipate claim 3. Claims 14 and 23 recite subject matter

addressed above with respect to claim 3 and are allowable for similar reasons.

IA(3). 35 U.S.C. § 102, Alleged Anticipation of Claims 4, 15, and 24

With respect to claim 4, the Examiner's Answer states:

Claims 4, 15 and 24 are dependent on dependent claims 3, 14 and 23 respectively. All three claims are basically claiming the same. Claim 4 stated as "the method as recited in claim 3, further comprising: determining whether data within the related set of data are at the same location as a service within the set of remaining services and responsive to data within the related set of data at the same location as a service within the set of remaining services, attaching the data to the service" as in the event that the video clip is only stored at this location, then a delay will be encountered as the DSI 30 waits for the video information to be downloaded. The PIM 22 will also direct that the number of FDVs to be decremented for this affected extended SRU 26 and the SRU under-run counter parameter identifies the location of over-accessed SRUs, audio-visual data will be moved or copied from heavily loaded SRUs to lightly loaded SRUs in an effort to distribute or flatten SRU demand (Fig. 1, col. 14, lines 28-38).

Appellant argues on page 20 as "Appellants submit that Kenner, in fact, does not teach or fairly suggest determining whether data within the related set of data are at the same location as a service within the set of remaining services." Again, Examiner respectfully disagrees with the Appellant because, Kenner clearly teaches two fold approach as in the event that the video clip is only stored at this location, then a delay will be encountered as the DSI 30 waits for the video information to be downloaded. The PIM 22 will also direct that the number of FDVs to be decremented for this affected extended SRU 26 and the SRU under-run counter parameter identifies the location of over-accessed SRUs, audio-visual data will be moved or copied from heavily loaded SRUs to lightly loaded SRUs in an effort to distribute or flatten SRU demand (Fig. 1, col. 14, lines 28-38).

Examiner's Answer dated December 8, 2005, pp. 30-31.

Kenner does not teach the feature of "determining whether data within the related set of data are at the same location as a service within the set of remaining services". The Examiner alleges that this feature is found in column 14, lines 28-38 of *Kenner*, which is reproduced above.

As previously mentioned, the passage cited by the Examiner discloses that a count parameter is used to identify over-accessed SRUs in a video clip and storage retrieval system, and copies of the video clips may be distributed from more heavily loaded SRUs to lightly loaded SRUs. A data sequencing interface (DSI) is used to increment the SRU counter and

communicates this information to the primary index managers (PIM). If the count for that SRU exceeds a threshold, the PIM directs requests for video clips on that SRU to an alternate SRU. If a video clip is only stored at the SRU, the DSI downloads the video clip. However, *Kenner* does not teach determining whether data within the related set of data are at the same location as a service within the set of remaining services. The passage above merely teaches that if an individual video clip is only stored at the failed SRU, the DSI downloads the video clip, rather than determining whether data within the related set of data are at the same location as a service within the set of remaining services. Consequently, *Kenner* does not anticipate claim 4. Claims 15 and 24 recite subject matter addressed above with respect to claim 4 and are allowable for similar reasons.

IB. 35 U.S.C. § 102, Alleged Anticipation of Claims 6-9, 11, 17-20, and 26-29

With respect to claim 6, the Examiner's Answer states:

Appellant argues on page 22 as "The applied reference fails to teach or suggest each and every claimed limitation. Therefore, *Kenner* does not anticipate claim 6." Examiner respectfully disagrees with the Appellant because, management of data sets by data services and duplicate copies of video clips placing on SRUs geographically closer to the users most interested in those videos. In both cases the information needed is placed closer so that the user can access easily and quicker than getting from a central location through communication network. Further, *Kenner* teaches to balance the load, SRU under-run count parameter identifies the location of "over-accessed" SRUs, audio-visual data will be moved or copied from these heavily loaded SRUs to more lightly loaded SRUs based on their under-run levels, in order to distribute or flatten SRU demand (co. 14, lines 33-38). In the Final Office Action, Every limitation of claim 6 has been addressed by explaining the reference with the citation.

Examiner's Answer dated December 8, 2005, pp. 31-32.

Kenner does not teach the feature of "responsive to an additional service joining the distributed set of services, querying management of the data within the related sets of data". The Examiner alleges that this feature is found in column 14, lines 33-38 of *Kenner*, which is reproduced above.

The Examiner states that an SRU under-run count parameter is used to identify over-accessed SRUs and allows data to be moved or copied from heavily loaded SRUs to more lightly loaded SRUs. However, the SRU under-run count parameter is not equivalent to a data service

being added to a distributed set of data services. Therefore, *Kenner* does not anticipate claim 6.

Independent claims 11, 17, and 26, as well as dependent claims 5, 16, and 25, recite subject matter addressed above with respect to claim 6 and are allowable for similar reasons. Since claims 7-9, 18-20, and 27-29 depend from claims 6, 17, and 26, the same distinctions between *Kenner* and the invention recited in claims 6, 17, and 26 apply for these claims. Additionally, claims 7-9, 18-20, and 27-29 recite other additional combinations of features not taught by the reference.

Therefore, Appellants respectfully request that the rejection of claims 6-9, 11, 17-20, and 26-29 under 35 U.S.C. § 102(e) not be sustained.

IB(1). 35 U.S.C. § 102, Alleged Anticipation of Claims 7, 18, and 27

With respect to claim 7, the Examiner alleges that the feature of “wherein the optimization criteria is based on location of the service within the distributed set of services” is found in Figure 1 and column 8, lines 38-47 and lines 28-30 of *Kenner*, which are reproduced above. As previously mentioned, the passages disclose identifying video clips by how often they are accessed (“accessed most frequently”), and that duplication of the identified most frequently accessed video clips are duplicated on SRUs that are closest to users that are accessing the video clips. Thus, *Kenner* teaches that the criteria for duplicating an individual video clip on another SRU is based on how often the video clip is accessed, rather than based on the location of the services. Consequently, *Kenner*, in fact, does not teach assigning management of data sets by data services based on location of the data services. Claims 18 and 27 recite subject matter addressed above with respect to claim 7 and are allowable for similar reasons.

IB(2). 35 U.S.C. § 102, Alleged Anticipation of Claims 8, 19, and 28

With respect to claim 8, the Examiner alleges that the feature of “examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service” is found in Figure 1 and column 14, lines 20-28 of *Kenner*, which are reproduced above. *Kenner* merely teaches that a data sequencing interface (DSI) is used to increment the SRU counter and communicates this information to the primary index managers (PIM), which directs requests for video clips on that SRU to an alternate SRU if the count for that SRU exceeds a threshold. Yet directing a request for an individual video clip to another SRU based on


the SRU counter is not the same as the related set of data managed by a failed service examined by the set of remaining services in the system. There is no mention in *Kenner* that remaining data services examine the related data sets managed by a failed data service. Consequently, *Kenner* does not anticipate claim 8. Claims 19 and 28 recite subject matter addressed above with respect to claim 8 and are allowable for similar reasons.

IB(3). 35 U.S.C. § 102, Alleged Anticipation of Claims 9, 20, and 29

With respect to claim 9, the Examiner alleges that the feature of “determining whether data within the related set of data are at the same location as a service within the set of remaining services” is found in Figure 1 and column 14, lines 28-38 of *Kenner*, which are reproduced above. However, *Kenner* makes no mention of determining whether data within the related set of data are at the same location as a service within the set of remaining services. *Kenner* merely teaches that if an individual video clip is only stored at the failed SRU, the DSI downloads the video clip, rather than determining whether data within the related set of data are at the same location as a service within the set of remaining services. Consequently, *Kenner* does not anticipate claim 4. Claims 9, 20, and 29 recite subject matter addressed above with respect to claim 4 and are allowable for similar reasons.

CONCLUSION

In view of the above, Appellants respectfully submit that claims 1-29 are allowable over the cited prior art and that the application is in condition for allowance. Accordingly, Appellants respectfully request the Board of Patent Appeals and Interferences to not sustain the rejections set forth in the Final Office Action.


Cathrine K. Kinslow
Reg. No. 51,886
YEE & ASSOCIATES, P.C.
PO Box 802333
Dallas, TX 75380
(972) 385-8777

CLAIMS APPENDIX

The text of the claims involved in the appeal reads:

1. A method of managing a set of data by a distributed set of services, comprising the steps of:

organizing the set of data into a plurality of related sets of data;

assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria; and

responsive to failure of a service within the distributed set of services, transferring management of the related set of data managed by the failed service to another service within the distributed set of services.

2. The method as recited in claim 1, wherein the optimization criteria is based on location of the service within the distributed set of services.

3. The method as recited in claim 1, further comprising:

detecting the failed service by a set of remaining services within the distributed set of services; and

examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service.

4. The method as recited in claim 3, further comprising:

determining whether data within the related set of data are at the same location as a service within the set of remaining services; and

responsive to data within the related set of data at the same location as a service within

the set of remaining services, attaching the data to the service.

5. The method as recited in claim 1, further comprising:

responsive to an additional service joining the distributed set of services, querying management of the data within the related sets of data; and
assigning management of a related set of data to the additional service within the distributed set of services based on the optimization criteria.

6. A method of managing a set of data by a distributed set of services, comprising the steps of:

organizing the set of data into a plurality of related sets of data;
assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria;
responsive to an additional service joining the distributed set of services, querying management of the data within the related sets of data; and
assigning management of a related set of data to the additional service within the distributed set of services based on the optimization criteria.

7. The method as recited in claim 6, wherein the optimization criteria is based on location of the service within the distributed set of services.

8. The method as recited in claim 6, further comprising:

detecting a failed service in the distributed set of services by a set of remaining services within the distributed set of services; and

examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service.

9. The method as recited in claim 8, further comprising:

determining whether data within the related set of data are at the same location as a service within the set of remaining services; and

responsive to data within the related set of data at the same location as a service within the set of remaining services, attaching the data to the service.

10. A data processing system, comprising:

a system bus;

a memory, including a set of instructions, functionally connected to the system bus; and

a processing unit functionally connected to the system bus, wherein the processing unit executes the set of instructions from the memory to organize a set of data into a plurality of related sets of data, wherein the data in each related set of data has at least one attribute between members, the processing unit assigns, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria, and, responsive to a failed service within the distributed set of services, the processing unit transfers management of the related set of data managed by the failed service to another service within the distributed set of services.

11. A data processing system, comprising:

a system bus;

a memory, including a set of instructions, functionally connected to the system bus; and

a processing unit functionally connected to the system bus, wherein the processing unit executes the set of instructions from the memory to organize a set of data into a plurality of related sets of data, wherein the data in each related set of data has at least one attribute between members, the processing unit assigns, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria, responsive to an additional service joining the distributed set of services, the processing unit queries management of the data within the related sets of data, and the processing unit assigns management of a related set of data to the additional service within the distributed set of services based on the optimization criteria.

12. A data processing system for managing a set of data by a distributed set of services, comprising:

organizing means for organizing the set of data into a plurality of related sets of data, wherein the data in each related set of data has at least one attribute between members;

assigning means for assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria; and

transferring means, responsive to a failed service within the distributed set of services, for transferring

management of the related set of data managed by the

failed service to another service within the distributed set of services.

13. The data processing system as recited in claim 12, wherein the optimization criteria is based on location of the service within the distributed set of services.

14. The data processing system as recited in claim 12, further comprising:

detecting means for detecting the failed service by a set of remaining services within the distributed set of services; and

examining means for examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service.

15. The data processing system as recited in claim 14, further comprising:

determining means for determining whether data within the related set of data are at the same location as a service within the set of remaining services; and

attaching means, responsive to data within the related set of data at the same location as a service within the set of remaining services, for attaching the data to the services.

16. The data processing system as recited in claim 12, further comprising:

querying means, responsive to an additional service joining the distributed set of services, for querying management of the data within the related sets of data; and

assigning means for assigning management of a related set of data to the additional service within the distributed set of services based on the optimization criteria.

17. A data processing system for managing a set of data by a distributed set of services, comprising:

organizing means for organizing the set of data into a plurality of related sets of data;

assigning means for assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria;

querying means, responsive to an additional service joining the distributed set of services, for querying management of the data within the related sets of data; and

assigning means for assigning management of a related set of data to the additional service within the distributed set of services based on the optimization criteria.

18. The data processing system as recited in claim 17, wherein the optimization criteria is based on location of the service within the distributed set of services.

19. The data processing system as recited in claim 17, further comprising:

detecting means for detecting a failed service in the distributed set of services by a set of remaining services within the distributed set of services; and

examining means for examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service.

20. The data processing system as recited in claim 19, further comprising:

determining means for determining whether data within the related set of data are at the same location as a service within the set of remaining services; and

attaching means, responsive to data within the related set of data at the same location as a service within the set of remaining service, attaching the data to the service.

21. A computer program product in a computer readable medium for managing a set of data by a distributed set of services, comprising:

instructions for organizing the set of data into a plurality of related sets of data;

instructions for assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria; and

instructions, responsive to a failed service within the distributed set of services, for transferring management of the related set of data managed by the failed service to another service within the distributed set of services.

22. The computer program product as recited in claim 21, wherein the optimization criteria is based on location of the service within the distributed set of services.

23. The computer program product as recited in claim 21, further comprising:

instructions for detecting the failed service by a set of remaining service within the distributed set of services; and

instructions for examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service.

24. The computer program product as recited in claim 23, further comprising:

instructions for determining whether data within the related set of data are at the same location as a service within the set of remaining services; and

instructions, responsive to data within the related set of data at the same location as a service within the set of remaining services, for attaching the data to the service.

25. The computer program product as recited in claim 21, further comprising:

instructions, responsive to an additional service joining the distributed set of service, for querying management of the data within the related sets of data; and

instructions for assigning management of a related set of data to the additional service within the distributed set of services based on the optimization criteria.

26. A computer program product in a computer readable medium for managing a set of data by a distributed set of services, comprising the steps of:

instructions for organizing the set of data into a plurality of related sets of data;

instructions for assigning, by a set of services, management of a related set of data to a service within the distributed set of services based on an optimization criteria;

instructions, responsive to an additional service joining the distributed set of services, for querying management of the data within the related sets of data; and

instructions for assigning management of a related set of data to the additional service within the distributed set of services based on the optimization criteria.

27. The computer program product as recited in claim 26, wherein the optimization criteria is based on location of the service within the distributed set of services.

28. The computer program product as recited in claim 26, further comprising:

instructions for detecting a failed service in the distributed set of services by a set of remaining services within the distributed set of services; and

instructions for examining, by the set of remaining services within the distributed set of services, the related set of data managed by the failed service.

29. The computer program product as recited in claim 28, further comprising:

instructions for determining whether data within the related set of data are at the same location as a service within the set of remaining services; and

instructions, responsive to data within the related set of data at the same location as a service within the set of remaining services, for attaching the data to the service.

EVIDENCE APPENDIX

There is no evidence to be presented.

RELATED PROCEEDINGS APPENDIX

There are no related proceedings.